
Subject: Re: [ckrm-tech] [RFC][PATCH][0/4] Memory controller (RSS Control)

Posted by [Paul Menage](#) on Mon, 19 Feb 2007 09:50:22 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 2/19/07, Kirill Korotaev <dev@sw.ru> wrote:

> >

> > I think it's OK for a container to consume lots of system time during
> > reclaim, as long as we can account that time to the container involved
> > (i.e. if it's done during direct reclaim rather than by something like
> > kswapd).

> hmm, is it ok to scan 100Gb of RAM for 10MB RAM container?

> in UBC patch set we used page beancounters to track container pages.

> This allows to make efficient scan controller and reclamation.

I don't mean that we shouldn't go for the most efficient method that's practical. If we can do reclaim without spinning across so much of the LRU, then that's obviously better.

But if the best approach in the general case results in a process in the container spending lots of CPU time trying to do the reclaim, that's probably OK as long as we can account for that time and (once we have a CPU controller) throttle back the container in that case. So then, a container can only hurt itself by thrashing/reclaiming, rather than hurting other containers. (LRU churn notwithstanding ...)

Paul
